

REMARKS

This is in response to the Office Action mailed 3/21/2007. It should be noted that all claims are commonly held jointly by between all joint inventors. Minor amendments have been made to claims 1-6, 8, 11, 14-15, 17-19, 21, 22-24, 28, 32-33, 35-38, and 45-47 for clarification purposes only. Also, claims 9-10, 12-13, 16, 20, 25-26, 29-31, 34, 39, 41-44 are hereby cancelled via the current amendment. Such amendments have been made without adding new matter. This response should obviate outstanding issues and make the pending claims allowable. Reconsideration of this application is respectfully requested in view of this response/amendment.

STATUS OF CLAIMS

1. Claims 1-47 are pending.
2. Claims 9-10, 12-13, 16, 20, 25-26, 29-31, 34, 39, 41-44 are hereby cancelled via the current amendment.
3. Claims 31 and 44 are rejected under 35 U.S.C. §112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention.
4. Claims 2-6, 8-21, 23-26, 28-39, and 41-47 are objected to because of minor informalities.
5. Claims 22-26 are rejected under 35 U.S.C. §101 because they are directed to non-statutory subject matter.
6. Claims 1-8, 15-19, 22-24, 26-28, 32, 38, 39, 40 and 47 are rejected under 35 U.S.C. §102(e) as being anticipated by Sitaraman et al. (US 2004/013627).

7. Claims 9, 10, 29, 30 and 41 are rejected under 35 U.S.C. §103(a) as being unpatentable over Sitaramann et al. (US 2004/013627).
8. Claims 12 and 25 are rejected under 35 U.S.C §103(a) as being unpatentable over Sitaraman et al. (US 2004/013627) in view of Veres et al. (6,807,156).
9. Claims 13, 14, 17, 21, 31, 35, 36, 37, 44, 45, and 46 are rejected under 35 U.S.C. §103(a) as being unpatentable over Sitaraman et al. (US 2004/013627) in view of Ni (US 6,298,055).
10. Claims 19, 20, 33 and 34 are rejected under 35 U.S.C. §103(a) as being unpatentable over Sitaraman et al. (US 2004/013627).

OVERVIEW OF CLAIMED INVENTION

A packetized streaming media delivery network carries many “streams” of differing media content. They often are from multiple sources and of different media types. The invention consists of a scalable hardware and/or software computing element resolving the network traffic into its individual streams for focused, simultaneous, and continuous real-time monitoring and analysis. The monitoring and analysis consists of delay factor and media loss rate which measure the cumulative jitter of the streaming media within the delivery network and the condition of the media payload. These measurements form a powerful picture of network problem awareness and resolution. The delay factor objectively indicates the contribution of the network devices in the streams’ path, allowing for both problem prediction and indication. In one example, tapping a packetized network at various locations allows for correlation of the same-stream performance at various network points to pinpoint the source(s) of the impairment(s).

REJECTIONS UNDER 35 U.S.C. §112

Claims 31 and 44 are rejected under 35 U.S.C. §112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention. The Examiner is respectfully requested to remove the 35 U.S.C. §112, second paragraph, rejection to claims 31 and 44 in light of their cancellation via the current amendment.

CLAIM OBJECTIONS

Claims 2-6, 8-21, 23-26, 28-39, and 41-47 are objected to because of minor informalities. The Examiner is respectfully requested to remove these objections to claims 2-6, 8-21, 23-26, 28-39, and 41-47 in light of the current amendment which is based on the Examiner's suggestions outlined on page 4 of the Office Action of 03/21/2007.

REJECTIONS UNDER 35 U.S.C. §101

Claims 22-26 are rejected under 35 U.S.C. 101 because they are directed to non-statutory subject matter. The Examiner is respectfully requested to remove the 35 U.S.C. §101 rejection to claims 22-26 in light of the current amendment which is based on the Examiner's suggestions outlined on page 5 of the Office Action of 03/21/2007.

REJECTIONS UNDER 35 U.S.C. §102

Claims 1-8, 15-19, 22-24, 26-28, 32, 38, 39, 40 and 47 are rejected under 35 U.S.C. 102(e) as being anticipated by Sitaraman et al. (US 2004/013627), hereafter Sitaraman. The rejection with respect to claims 16, 20, 26, and 39 are moot in view of their cancellation via

the current amendment. Applicants respectfully disagree with the Examiner that the claims are taught by the cited art. The Manual for Patenting Examining Procedure (MPEP) § 2131 clearly sets forth the standard for rejecting a claim under 35 U.S.C. § 102(b). “A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference.” (MPEP § 2131, quoting *Verdegaal Bros. v. Union Oil Co. of California* 2 USPQ2d 1051, 1053 (Fed Cir. 1987)). In this case, the cited art (i.e., Sitaraman) fails to teach the claimed invention as required by the MPEP.

There are several significant differences between the present invention and the Sitaraman reference. For instance, the fundamental goal of the measurements (monitoring) is different. The present invention characterizes the overall impact to media stream(s) (by network devices and media sources) in a highly scalable way, by filtering to isolate the streams and analyze each individually and simultaneously. Sitaraman treats the media consumer (i.e. decoder) as a black box and attaches measurements probes to both sides of the decoder, input being the network and output being the renderer. Sitaraman’s approach, although still targeting fundamentally different measurements of quality/performance, requires one decoder and renderer (media consumer) for each stream it wishes to monitor, if streams are of differing types this would include many different types of media consumers. Their approach is neither practical nor cost effective to scale to a network that is carrying hundreds or thousands of media streams simultaneously.

The presently claimed invention does not require any interaction with the media consumer nor media source, providing a scalable solution which is quite novel. It should be

emphasized that prior to the present invention's solution, it has not been possible to simultaneously and cost effectively analyze/monitor a multiplicity of streams individually.

With respect to claim 1, it is worth emphasizing that the present invention does not simply isolate one stream to resolve, but resolves all streams all the time for continual and simultaneous monitoring/analysis. Such a teaching is neither explicit nor implicit in the teachings of Sitaraman.

With respect to claim 2, the Examiner argues that Sitaraman's use of "rebuffer-time-per-minute" corresponds to a delay factor. Sitaraman's use of rebuffer-time-per-minute is a measurement of how much time is spent in a stalled "waiting for rebuffer" state per minute. The present invention's "Delay Factor" (DF) is different and unique. The present invention's DF represents the "time value" of any buffering (regardless of whether or not a media consumer is stalled for rebuffering... there is no relationship between the "time value" of buffering and the act of rebuffering due to buffers running empty). DF represents the "time" in which you would need to delay the media by buffering to contain the media stream as well as absorb observed network jitter accumulation on the network so that the media consumers' buffers neither overflow nor empty. The simple act of packetizing media requires some level of buffering (i.e. delay) simply to hold the packet's media payload when it arrives.

In one non-limiting example, DF is a calculated buffer size divided by the bitrate of the media. For instance, assume one were to buffer 10 packets of a stream of bitrate "A" such that the delay was 4ms (those 10 packets took 4ms to arrive given bitrate "A"), thus the DF=4ms. If

one were to then buffer 10 packets of a different stream of bitrate “1/2-A”, it would take twice as long for those 10 packets to arrive, thus the DF=8ms. Thus, no relation to a per-unit-time measurement of a media consumers’ time spent in a “rebuffering” state.

Hence, Sitaraman fails to teach or suggest the present invention’s feature of a delay factor.

With respect to claim 3, the Examiner appears to equate Sitaraman’s “Packet Loss Rate” to the present invention’s Media Loss Rate. It should be noted that although the terms sound similar, they signify distinct meanings. The distinction here is that MLR references the loss rate of the media itself; i.e., the loss detected within the payload of the network packets as opposed to the loss of the packets themselves. Many media data formats (i.e. the data carried as payload in the network packets) include enough information to determine if there is loss or corruption in the media itself, making the distinction between loss rate of the media (MLR) and loss rate of the network packets. It is both possible and common that one will observe media loss (MLR) without simultaneously observing network packet loss. Typical packetized streaming media networks have segments that may transmit the streaming media natively, without packetization, then (re)packetize at the end of that segment. This allows for the media data itself to be corrupted without experiencing any network packet loss or detectable impairment. This further extends to the generator of the media data having been in error prior to packetization for transmission on the packetized network.

It is worth emphasizing that the term “streaming media” with regards to the present invention refers to a continuous flow of media data from a media source to a media destination intended to be delivered just-in-time as needed at the destination without underflow or overflow in order to allow the destination to use the media data in its partiality. The term “streaming media over a packetized network” with regards to the present invention refers to streaming media data being encapsulated into the payload of native packets appropriate for the network in question. This is generally a bunching up of some amount of streaming media data into the packetized network packet’s payload. The term “stream” with regards to the present invention refers to a specific uniquely identifiable packetized streaming media channel contained on a network. A stream does not imply all aggregate traffic on the network. For example, consider the open-air carrying multiple television channels. A stream would be one specific television channel. The term “media consumer” or “data consumer” with regards to the present invention refers to the intended target(s) of the media stream in question. For example, if the media is video, the media or data consumer would be the decoder(s) that the media is being streamed to.

The above-mentioned arguments for claim 1 substantially apply to claim 4. Further, “minimal distortion” is clarified as the determination of the arrival time of a given packet as close to when it was actually received at the network interface. This should be distinguished from buffering wherein a packet is held in a buffer for some variable length of time before assigning it an arrival time (timestamp).

Hence, at least with respect to the above-presented arguments, Sitaraman fails to render obvious the teachings of claim 4.

The above presented arguments with respect to dependent claims 2 and 3 substantially apply to dependent claims 5 and 6 as they recite similar features.

Further, the above-presented arguments regarding independent claims 1 and 4 substantially apply to independent claims 7, 22, 27, and 40.

With respect to the rejections of claims 38 and 47, the Examiner appears to equate the “central management” feature to the “forwarding a stream media payload” aspect of Applicants claimed invention. The features of claims 38 and 47 are quite different from forwarding “stats” to a stats/data consumer (central management). Claims 38 and 47 deal with stripping out the media content itself (“forwarding a streaming media *payload*”) from the network packet streams and forwarding to a media type-dependent interface in order to connect to 3rd party equipment designed to deal with that media type when not on a packetized delivery network. Let’s take the case of MPEG2-TS where there are numerous 3rd party vendors that make equipment to deal with (decode, analyze, store, etc) an MPEG2-TS stream on a media type-dependent interface, DVB-ASI for example. However, until the present invention’s solution, there was no solution for isolating a media stream on a packetized network (non media type-dependent, such as Ethernet) and ripping out just the media payload itself and forwarding to such a media type-dependent interface (such as DVB-ASI) such that 3rd party vendors’ equipment could gain access to that media stream data.

With regard to the objection of claim 11, the Examiner appears to equate the “Startup time” mentioned in Sitaraman with the “control information” aspect of claim 11. “Startup time” is an observed measurement akin to the time in which the stream was first observed rather than control information. Control information, on the other hand, is comprised of instructions that may be sent from a media source (such as video encoder) to a media consumer (such as video decoder) or visa versa. This would include such traffic as a user requesting that a given media stream speed up, slow down, jump backwards, etc. This is deemed “control plane” traffic on the packetized network. For reference, the streaming media would be considered “data plane” traffic on the packetized network. Sitaraman fails to teach or suggest such a feature.

With respect to Claim 15, it should be noted that the flow rate differs from bitrate in that bitrate refers to the media stream payload whereas flow rate includes the overhead of packetization.

With regards to the rejection of claim 18, it should once again be re-emphasized that the “rebuffer-time-per-minute” is not related to Applicants’ delay factor or DF. In claim 18, a flow rate balance involves determining the so called “balance” point of a buffer given a particular flow rate. If the packet arrival rate into a buffer exactly matched the media consumption rate at the output of the buffer, the balance point would be the point of buffer fullness with a single network packet’s worth of media payload. With the rates being equal, the buffer would be in “balance” and would neither trend toward empty nor trend toward full. The buffer fullness would approach empty at the same instant another packet arrives, thus the IFRB wanders between a packet’s worth of media payload in the buffer to 0, continually repeating (never filling

more and never becoming truly empty). To tie into DF, the total “wander” of the IFRB represents the minimum amount of buffering required (so called Virtual Buffer), whereas this buffer size divided by the media bitrate equals DF. The Instantaneous Flow Rate Balance (IFRB) is a snap-shot of the buffer fullness level. Instantaneous Flow Rate Balance Deviation would be how far the IFRB is displaced from nominal “balance.” This is a measurement of whether or not the buffer is trending toward empty or trending toward full due to stream source or network devices (i.e. network effects). Claim 18 refers to a specific implementation of measuring IFRB Deviation. The example was made purposely simplistic so as to be easily followed, a brute force “counter” approach. The IFRB (and ultimately DF) calculation involves tracking the bitrate of the media payload in isolated stream(s) according to the media type(s) and relating this directly to the pattern of arrival times (via timestamp) observed for the stream(s) each individually and simultaneously.

While a “counter” by itself is generic, what it counts is specific. In our case, we are using a counter to refer to a calculation element coupled with a memory storage element. Further, the present invention’s counter does not simply counting - 1, 2, 3, 4, etc. We calculate the IFRB, store it, compare it to the past calculated IFRB for this stream, store the max seen, store the min seen, etc.

Such features are neither taught for nor suggested by the Sitaraman reference.

Further, the examiner appears to draw a parallel between IFRB/DF and Sitaraman’s measurement of “rebuffers-per-minute”, stating that rebuffering correlates to buffer size which

then correlates to network jitter. This is untrue in its entirety. While “rebuffering” does relate to buffer size and jitter under certain circumstances it fails to represent the same fundamental that IFRB/DF captures. “Rebuffering” does not capture the “accumulation” of jitter on the network for a stream. Various network devices in the distribution path of stream(s) affect the overall jitter profile and the jitter accumulation. Jitter ebbs and flows (accumulates and dissipates) as the various network devices contain varying amount of packet handling resources and various QoS (Quality of Service) algorithms. This places additional “stress” and requirement on network devices downstream, including the final media consumer. Buffer underflows and overflows occur at network devices as well as the final media consumer.

Measuring “rebuffering” as per US2004/0136327 fails to: (a) capture/distinguish buffer overflows; (b) predict impending buffer overflow or underflow, as a rebuffer event indicates a bad thing has already happened; (c) yield buffer size recommendations for an equipment designer or network operator in order to support the observed network jitter accumulation; and (d) describe the “burstiness” (packets bunching up) of a stream, which is an important factor in streaming media network implementations and the network devices used in packetized streaming media network implementations.

The presently claimed invention, on the other hand, remedies the failure of the prior art.

REJECTIONS UNDER 35 U.S.C. §103

Claims 13, 14, 17, 21, 31, 35, 36, 37, 44, 45, and 46 are rejected under 35 U.S.C. 103(a) as being unpatentable over Sitaraman et al. (US 2004/013627) in view of Ni (US 6,298,055).

The rejection with respect to claims 13, 31, and 44 are considered moot in view of their rejection via the current amendment. To be properly rejected under 35 U.S.C. 103(a), the combination of cited references should disclose all of the features of the rejected claims.

The above mentioned arguments with respect to the Sitaraman reference's applicability to the independent claims 1, 4, 7, 22, 27, and 40 substantially applies to pending dependent claims 14, 17, 21, 35, 36, 37, 45, and 46 as they inherit all the features of the claim from which they depend from. Hence, Applicants respectfully assert that the combination of Sitaraman and Ni cannot teach each and every feature of claims 14, 17, 21, 35, 36, 37, 45, and 46.

Further, Ni deals with "flow control" packets that are designed to control the traffic flow pattern on the network. As a network approaches 100% capacity, there exists several congestion points within a network. Proper flow control of a full-duplex network architecture requires reliable delivery of the flow control packets from receiver to transmitter. Should a flow control packet be lost, the data flow that was being controlled may likely experience loss. In the event the lost control packet was saying "stop sending," the transmitter would not stop causing overflow loss at the receiver. In the event the lost control packet was saying "begin sending again," the transmitter would not begin, causing the buffers at receiver to run empty and/or the buffers at the transmitter to fill causing either overflow loss here or pushing back, via its own flow control packets upstream, thereby degrading performance. Ni correctly states that when the network is lightly loaded, there is a greatly reduced chance of a flow control packet from receiver to transmitter being lost in the network, which in turn implies reduced chance of overflow loss and/or performance degradation. Ni continues to state that when the network is

heavily loaded, the converse is true, it is very likely a flow control packet will be lost from receiver to transmitter greatly increasing the chances of overflow loss and/or performance degradation.

The core of Ni's argument centers about scaling the frequency of transmission (i.e. bandwidth required) of flow control packets from receiver to transmitter according to the network loading level. Continuing to assert that in the case of high network loading when flow control packets experience greater chance of being lost, more frequent flow control packets will be sent to compensate for loss. In the lightly loaded network case, the bandwidth required for flow control packets can be reduced by sending them less frequently since there is much less chance of lost flow control packets.

The fundamental application and targeted invention of Sitaraman and Ni differ greatly from the presently claimed invention. Also, it is worth mentioning that with regard to the reporting of calculated statistics, Sitaraman does not make any mention of method of forwarding the stats nor does it discuss any adjustment to stats calculation frequency based upon available bandwidth to report such calculations. Ni allocates control traffic bandwidth via a QoS algorithm for the expressed purpose of modifying and altering the network traffic itself. It does not scale the bandwidth of the flow control packets in order to make available additional network capacity, it does so to increase the likelihood that the flow control packets are received by the transmitter in a timely manner.

Furthermore, as network available capacity decreases (network heavily loaded), the approach of Ni is to further reduce the network capacity by increasing the bandwidth consumed by the flow control packets. These claims of our invention intend to scale the bandwidth consumed by the stats reporting based upon several factors, including a QoS algorithm or the total aggregate bandwidth available with a given network technology. Scaling the bandwidth consumed by stats reporting is done to allow more available network capacity, a stark contrast from the teachings of Ni.

Further, there is no control exhibited over or alteration/modification of the streams' flow as a result of our claims. The fundamental reasons between our claims and Ni's for scaling the bandwidth are considered unrelated. Prior solutions for increasing network available bandwidth in the presence of any stats reporting have been to purchase additional network devices to increase capacity as opposed to scaling bandwidth consumption of stats reporting itself. Our claims do not represent an obvious extension to Ni as persons having ordinary skill in the art of controlling the flows of full-duplex networks would not conceive of our approach as indicated in the claims.

Claims 37 and 46 teach statically scaling the stats computation frequency and/or number of stats calculated in proportion to the reporting network technology's aggregate bandwidth, regardless of network loading. For instance, there is 1000x more available aggregate bandwidth on a 10 Gb/s network vs. a 10 Mb/s network. As such, one would likely wish to compute fewer stats and/or less frequently when the reporting network is 10 Mb/s capable as compared to 10 Gb/s capable. It should be noted that the scaling is done once (static), rather than

continuously (dynamic). Such static scaling is done at the time in which the reporting network is chosen and is changed only if the reporting network choice has changed. Such features are not rendered obvious by the combination of the Sitaraman and Ni references.

With regards to the rejection of claims 19 and 33, the Examiner states that Sitaraman indicates that counter overflow would be avoided by periodic clearing. Applicants are unable to find any such reference in the Sitaraman reference. As per previous discussion, the “counter” was to reflect a generic simple case example, wherein the “counter” generally represented a calculation element and a storage element. Although a PLL/DLL may have been obvious as a means of syncing to a bitrate of a media stream, clearing all memory of an IFRB and virtual buffer as a means of syncing and removing drift or frequency offset is not considered obvious. It should be emphasized that where as a PLL/DLL would have removed drift and frequency offset, it does not periodically clear any storage memory.

SUMMARY

As has been detailed above, none of the references, cited or applied, provide for the specific claimed details of Applicants' presently claimed invention, nor renders them obvious. It is believed that this case is in condition for allowance and reconsideration thereof and early issuance is respectfully requested.

As this response has been timely filed, no request for extension of time or associated fee is required. However, the Commissioner is hereby authorized to charge any deficiencies in the fees provided to Deposit Account No. 50-4098.

If it is felt that an interview would expedite prosecution of this application, please do not hesitate to contact Applicants' representative at the below number.

Respectfully submitted,

/ramraj soundararajan/

Ramraj Soundararajan
Registration No. 53,832

IP Authority, LLC.
9435 Lorton Market Street
#801
Lorton, VA 22079
571-642-0033

June 21, 2007

APPENDIX A